# Sun Solaris

# HP StorageWorks Disk Array XP operating system configuration guide

XP48
XP128
XP512
XP1024
XP12000

fifth edition (August 2004)

part number: A5951-96033

This guide describes the requirements and procedures for connecting the XP family of disk arrays to a Sun system and configuring the new disk array for operation with Solaris.

Printed in the U.S.A.

*HP StorageWorks Disk Array XP Operating System Configuration Guide: Sun Solaris*

fifth edition (August 2004)
part number: A5951-96033

# Contents

About this guide 5
Intended audience 5
Disk arrays 5
Related documentation 5
Conventions 6
HP technical support 7
HP storage website 7
HP authorized reseller 7
Revision history 8
Warranty statement 9

**1 Installation 11**
Features and requirements 12
Fibre Channel interface 13
Device types 13
Failover 14
SNMP configuration 15
RAID Manager command devices 15
Installation procedures 16
Install and configure the disk array 17
Setting the System Option Modes 17
Configuring the Fibre Channel ports 18
Setting the Host Mode for the disk array ports 20

Install and configure the host 21
Loading the OS and software 21
Installing the HBAs 21
Configuring the HBAs 21
Configuring Sun HBAs 22
Configuring JNI HBAs 23
Configuring Emulex HBAs 27
Verifying the HBA configuration 29
Clustering and Fabric zoning 30
Fabric zoning and LUN security for multiple operating systems 31
Connect the disk array 32
Defining the paths 33
Adding the new device paths to the system 34
Verifying disk array device recognition 35
Configure disk array devices 36
Partitioning and labeling the devices 37
Creating the file systems 39
Configure for use with Veritas Volume Manager 40

**2 Troubleshooting 41**
Error conditions 42
Calling the HP support center 44

**A Disk array device emulations 45**
Supported emulations 46
Device type specifications 47
Device Geometry Tables 50
OPEN-x geometry table 50
OPEN-3*n LUSE device geometry table 51
OPEN-8*n LUSE device geometry table 53
OPEN-9*n LUSE device geometry table 55
OPEN-E*n LUSE devices geometry table 57
OPEN-L*n LUSE device geometry table 58
OPEN-x*n VLL-LUSE devices geometry table (example) 59
OPEN-V*n VLL-LUSE device geometry table (example) 60

**Glossary 61**

**Index 65**

# About this guide

This guide describes the requirements and procedures for connecting the XP family of disk arrays to a Sun host, and configuring the disk array for use with the Solaris operating system.

## Intended audience

This guide is intended for system administrators who have knowledge of the following topics:

- Data processing concepts
- Direct access storage device subsystems and their basic functions
- Disk arrays and RAID technology
- Operating system commands and utilities

## Disk arrays

Unless otherwise noted, the term *disk array* refers to these disk arrays:

HP Surestore Disk Array XP512
HP Surestore Disk Array XP48
HP StorageWorks Disk Array XP128
HP StorageWorks Disk Array XP1024
HP StorageWorks XP12000 Disk Array

## Related documentation

HP provides the following related documentation:

- *HP StorageWorks Disk Array XP128: Owner's Guide*
- *HP StorageWorks Disk Array XP1024: Owner's Guide*
- *HP StorageWorks XP12000 Disk Array: Owner's Guide*

For information about operating system commands and third-party products, refer to the manufacturer's documentation.

## Conventions

This guide uses the following text conventions.

| | |
|---|---|
| Figure 1 | Blue text represents a cross-reference. For the online version of this guide, the reference is linked to the target. |
| www.hp.com | Underlined, blue text represents a website on the Internet. For the online version of this guide, the reference is linked to the target. |
| **literal** | Bold text represents literal values that you type exactly as shown, as well as key and field names, menu items, buttons, file names, application names, and dialog box titles. |
| *variable* | Italic type indicates that you must supply a value. Italic type is also used for manual titles. |
| `input/output` | Monospace font denotes user input and system responses, such as output and messages. |
| *Example* | Denotes an example of input or output. The display shown in this guide may not match your configuration exactly. |
| [ ] | Indicates an optional parameter. |
| { } | Indicates that you must specify at least one of the listed options. |
| \| | Separates alternatives in a list of options. |

## HP technical support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week. Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support:

http://h18006.www1.hp.com/storage/arraysystems.html

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

## HP storage website

For the most current information about HP StorageWorks XP products, visit the support website. Select the appropriate product or solution from this website:

http://h18006.www1.hp.com/storage/arraysystems.html

For information about product availability, configuration, and connectivity, consult your HP account representative.

## HP authorized reseller

For the name of your nearest HP authorized reseller, you can obtain information by telephone:

| | |
|---|---|
| United States | 1-800-345-1518 |
| Canada | 1-800-263-5868 |
| Or contact: | www.hp.com |

# Revision history

| | |
|---|---|
| September 1999 | Open-8 emulation added. |
| January 2000 | Content extensively revised and reorganized. |
| June 2000 | Added support for XP512.<br>Content reorganized and revised. |
| February 2001 | Added appendices and glossary. |
| November 2003 | Changed brand name to StorageWorks.<br>Added support for OPEN-L and OPEN-V.<br>Changed CVS to VSC. |
| August 2004 | Updated for XP12000 and technical improvements. |

# Warranty statement

HP warrants that for a period of ninety calendar days from the date of purchase, as evidenced by a copy of the invoice, the media on which the Software is furnished (if any) will be free of defects in materials and workmanship under normal use.

**DISCLAIMER. EXCEPT FOR THE FOREGOING AND TO THE EXTENT ALLOWED BY LOCAL LAW, THIS SOFTWARE IS PROVIDED TO YOU "AS IS" WITHOUT WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED. HP SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE.** Some jurisdictions do not allow exclusions of implied warranties or conditions, so the above exclusion may not apply to you to the extent prohibited by such local laws. You may have other rights that vary from country to country, state to state, or province to province.

WARNING! **YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT USE OF THE SOFTWARE IS AT YOUR SOLE RISK.** HP DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, VIRUS-FREE OR ERROR-FREE, OR THAT DEFECTS IN THE SOFTWARE WILL BE CORRECTED. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. HP DOES NOT WARRANT OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE OR RELATED DOCUMENTATION IN TERMS OF THEIR CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY HP OR HP'S AUTHORIZED REPRESENTATIVES SHALL CREATE A WARRANTY.

**LIMITATION OF LIABILITY. EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT INCLUDING NEGLIGENCE WILL HP OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.** Your use of the Software is entirely at your own risk. Should the Software prove defective, you assume the entire cost of all service, repair or correction. Some jurisdictions do not allow the exclusion or limitation of liability for incidental or consequential damages, so the above limitation may not apply to you to the extent prohibited by such local laws.

**NOTE. EXCEPT TO THE EXTENT ALLOWED BY LOCAL LAW, THESE WARRANTY TERMS DO NOT EXCLUDE, RESTRICT OR MODIFY, AND ARE IN ADDITION TO, THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE LICENSE OF THE SOFTWARE TO YOU; PROVIDED, HOWEVER, THAT THE CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS IS SPECIFICALLY DISCLAIMED AND SHALL NOT GOVERN OR APPLY TO THE SOFTWARE PROVIDED IN CONNECTION WITH THIS WARRANTY STATEMENT.**

# 1

# Installation

Installation of the HP StorageWorks Disk Array XP is performed by your HP service representative and you. The HP service representative installs the disk array and formats the disk devices. You configure the host server for the new devices with assistance from the HP service representative.

# Features and requirements

The disk array and host have the following features and requirements.

- HP StorageWorks disk arrays:

  **XP48:** Up to 48 drives from 72 GB to 8.7 TB, 24 FC ports

  **XP128:** From 8 to 128 drives for up to 18 TB, 48 FC ports

  **XP512:** Up to 512 drives from 72 GB to 93 TB, 48 FC ports

  **XP1024:** From 8 to 1024 drives for up to 149 TB, 64 FC ports

  **XP12000:** Up to 1152 drives for up to 165 TB, 128 FC ports

- Sun server
- Sun Solaris operating system, version 8.0 or later with current patches
- Solaris Volume Manager or Veritas Volume Manager
- Latest SAN Foundation Kit with current patches

  Sun StorEdge Traffic Manager requires that you configure /kernel/drv/scsi_vhci.conf

  For SAN information, refer to http://www.sun.com/storage/san/ for the *Sun StorEdge SAN Foundation Software & Installation Guide* and the *Sun StorEdge Traffic Manager Software Installation and Configuration Guide.*

- Host Bus Adapters (HBAs): Install adapters and all utilities and drivers. Refer to the adapter documentation for installation details.
- *(Recommended)* HP StorageWorks Command View XP with LUN management feature or Remote Control with the LUN Configuration Manager XP option for configuring disk array ports and paths.
- *(Recommended)* HP StorageWorks Secure Manager XP: Allows the host to access only array devices for which it is authorized.
- Other available XP Software (some may not apply to your system):
  HP StorageWorks Business Copy XP
  HP StorageWorks Continuous Access XP
  HP StorageWorks Continuous Access Extension XP
  HP StorageWorks Auto LUN XP

HP StorageWorks Data Exchange XP
HP StorageWorks Resource Manager XP
HP StorageWorks RAID Manager XP
HP StorageWorks Cache LUN XP
HP StorageWorks Auto Path XP
HP StorageWorks Cluster Extension XP
HP StorageWorks Performance Advisor XP software

## Fibre Channel interface

The XP48, XP128, XP512, XP1024, and XP12000 disk arrays support these 1 Gbps and 2 Gbps Fibre Channel interfaces:

- Short-wave non-OFC (open fiber control) optical interface
- Multimode optical cables with SC or LC connectors
- Public or private arbitrated loop (FC-AL) or fabric direct attach
- Fibre Channel switches

Even though the interface is Fibre Channel, this guide uses the term “SCSI disk” because disk array devices are defined to the host as SCSI disks.

## Device types

The disk arrays support the following device types:

- **OPEN-*x* devices:** OPEN-x logical units represent disk devices. Except for OPEN-V, these devices are based on fixed sizes. OPEN-V is a user-defined size. Supported emulations include OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-L, and OPEN-V devices.
- **LUSE devices (OPEN-*x***n*):** Logical Unit Size Expansion (LUSE) allows you to combine 2 to 36 OPEN-x devices to create expanded LDEVs larger than standard OPEN-*x* disk devices. For example, an OPEN-x LUSE volume created from ten OPEN-x CVS volumes is designated as OPEN-x*10.
- **CVS devices (OPEN-x CVS):** Volume Size Configuration (VSC) defines custom volumes (CVS) that are smaller than normal

fixed-sized logical disk devices (volumes). (OPEN-V is a CVS-based custom disk size that you determine. OPEN-L does not support CVS.)

- **LUSE (expanded) CVS devices (OPEN-x*n CVS):** LUSE CVS combines CVS devices to create an expanded device. This is done by first creating CVS custom-sized devices and then using LUSE to combine from 2 to 36 CVS devices. For example, if three OPEN-9 CVS volumes are combined to create an expanded device, this device is designated as OPEN-9*3-CVS.

## Failover

The disk arrays support many standard software products that provide host, application, or I/O path failover and logical volume (storage) management.

These include the VERITAS Cluster Server and Sun Cluster host failover products for the Sun Solaris operating system. Refer to the documentation for these products and Sun technical support for installation and configuration information. Your HP representative may have to set specific disk array System Option Modes for these products.

### Multipath management and I/O path failover

Sun supplies software called Sun StorEdge Traffic Manager (STMS) and Veritas offers Veritas Volume Manager (VxVM), both of which provide multipath configuration management, HBA I/O load balancing, and both automatic and manual failover support. STMS requires Sun host bus adapters. Hitachi Data Link Manager (HDLM) is another failover software product available for Solaris.

Refer to the manufacturers' manuals for instructions. STMS and VxVM installation manuals can be downloaded from the Sun website at www.sun.com and the Veritas website at www.veritas.com, respectively.

## SNMP configuration

The disk arrays support standard Simple Network Management Protocol (SNMP) for remotely managing the disk array from the host. The SNMP agent on the remote console PC or Command View can provide status and Remote Service Information Message (R-SIM) reporting to the SNMP manager on the host for up to eight disk arrays. To configure the SNMP manager on the host, refer to the operating system documentation.

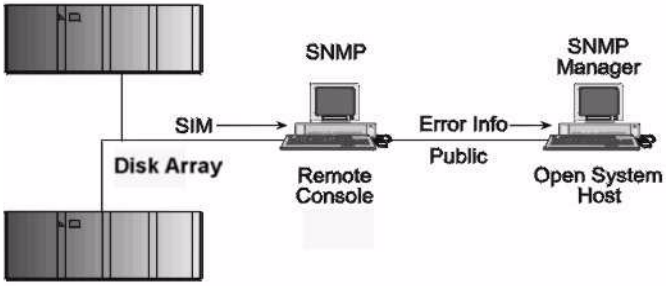


## RAID Manager command devices

RAID Manager manages Business Copy (BC) and/or Continuous Access (CA) operations from a server host. To use RAID Manager with BC or CA, you must use Command View or LUN Configuration Manager to designate at least one LDEV as a command device. Refer to the Command View or LUN Configuration Manager user guide for information about how to designate a command device.

# Installation procedures

The HP representative and you perform the following procedures:

1. "Install and configure the disk array" (page 17)
   - "Setting the System Option Modes"
   - "Configuring the Fibre Channel ports"
   - "Setting the Host Mode for the disk array ports"
2. "Install and configure the host" (page 21)
   - "Loading the OS and software"
   - "Installing the HBAs"
   - "Configuring the HBAs"
   - "Verifying the HBA configuration"
   - "Clustering and Fabric zoning"
   - "Fabric zoning and LUN security for multiple operating systems"
3. "Connect the disk array" (page 32)
   - "Defining the paths"
   - "Adding the new device paths to the system"
   - "Verifying disk array device recognition"
4. "Configure disk array devices" (page 36)
   - "Partitioning and labeling the devices"
   - "Creating the file systems"
5. "Configure for use with Veritas Volume Manager" (page 40)

# Install and configure the disk array

The HP service representative performs the following tasks:

- Assembling hardware and installing software
- Loading the microcode updates
- Installing the channel adapters (CHAs) and cabling
- Installing and formatting devices

You perform the additional tasks below. If you do not have Command View or LUN Configuration Manager, your HP service representative can perform these tasks for you.

## Setting the System Option Modes

The HP representative sets the System Option Mode(s) based on the operating system and software configuration of the host.

# Configuring the Fibre Channel ports

Configure the disk array Fibre Channel ports by using Command View (shown) or the Fibre Parameter window in LUN Configuration Manager. Select the settings for each port based on your storage area network topology. Use switch zoning if you connect different types of hosts to the array through the same switch.

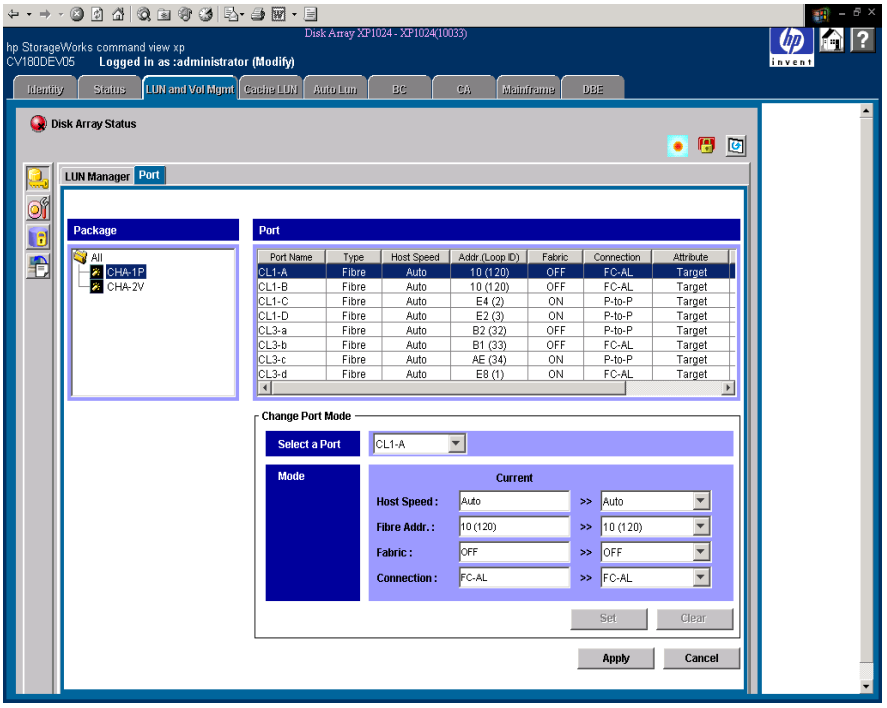


## Fibre Address

In fabric environments, the port addresses are assigned automatically. In arbitrated loop environments, you set the port addresses by selecting a unique arbitrated loop physical address (AL-PA) or loop ID for each port.

## Fabric and Connection parameter settings

You can set each array port to FABRIC ON or OFF with connections of POINT-TO-POINT or FC-AL as shown in the following table and figures. For detailed topology information, refer to the *HP StorageWorks SAN Design Reference Guide* on the hp.com website.

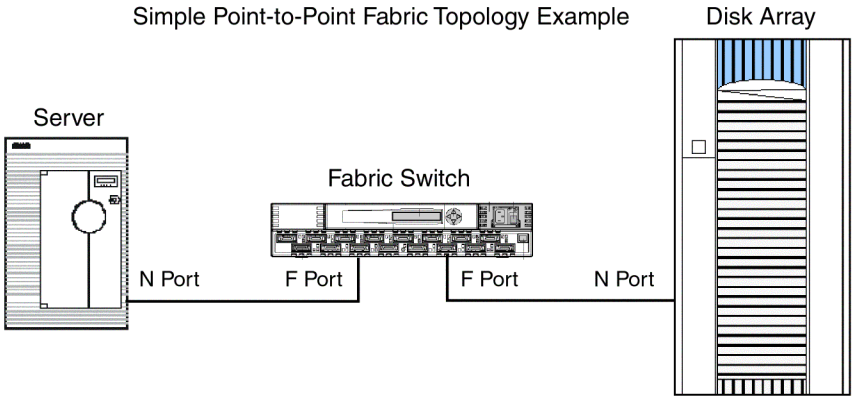


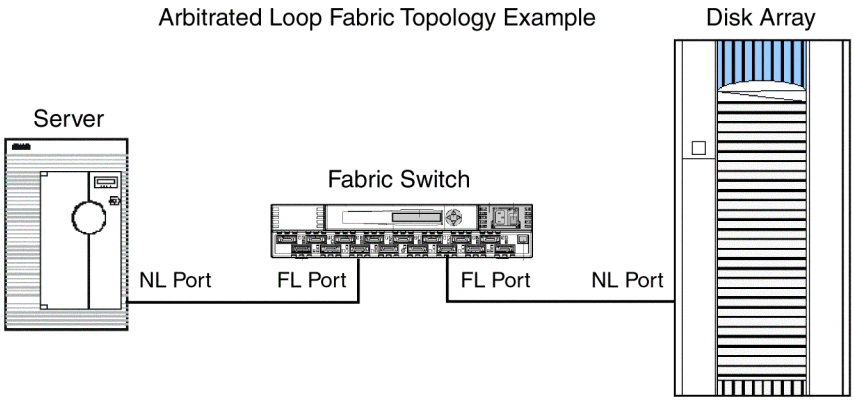


| Fabric Parameter | Connection Parameter | Provides |
|---|---|---|
| ON | FC-AL | NL-port (SAN public arbitrated loop) |
| ON | Point-to-Point | N-port (SAN fabric port) |
| OFF | FC-AL | NL-port (private arbitrated loop; direct connect without a SAN) |
| OFF | Point-to-Point | Not supported |

## Setting the Host Mode for the disk array ports

The disk array ports have Host Modes that you must set depending on the host you use. After the disk array is installed, use Command View (shown) or LUN Configuration Manager to set the Host Mode for each port.

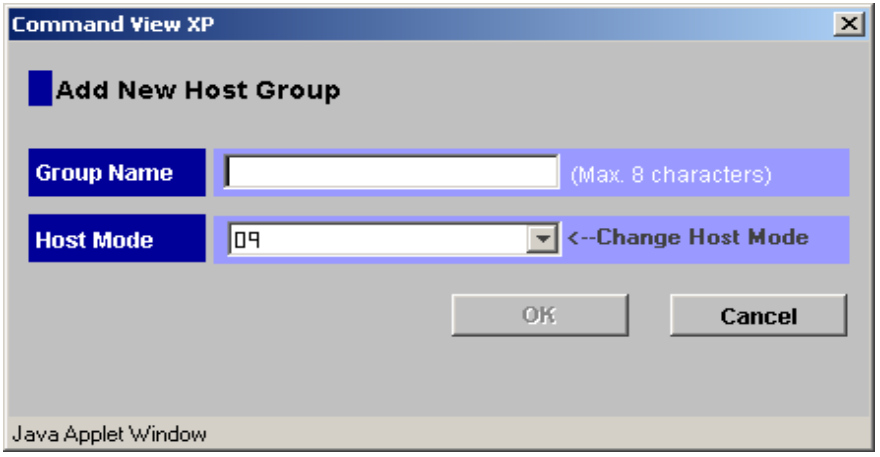


The required host mode setting for Solaris is **09**.

# Install and configure the host

Install and configure the host and host bus adapters (HBAs) that connect the host to the disk array.

## Loading the OS and software

Follow the manufacturer's instructions to load the operating system and software onto the host. Load all OS patches and configuration utilities supported by HP and the HBA manufacturer.

## Installing the HBAs

Install the host bus adapters using the HBA manufacturer's instructions.

## Configuring the HBAs

Load and configure the HBA driver software and setup utilities according to the manufacturer's instructions. Configuration settings specific to the XP array differ depending on the manufacturer. See the following pages for details about configuring HBAs from specific manufacturers.

### WWN

The HBA configuration process may require you to enter the WWN for the array port(s) to which it will connect. The HP representative can provide you this information, or you may display this information on the SAN switch.

## Configuring Sun HBAs

Sun HBAs have the following configuration requirements:

- The HP representative must set array System Option Mode 244 (ON).
- Install the latest SUN Foundation Kit with associated patches. Use the Sun supplied **install_it** script to automate installation. The script is available at http://www.sun.com/storage/san/.
- If you use Sun StorEdge Traffic Manager (MPxIO), edit the driver configuration file (**/kernel/drv/scsi_vhci.conf**) to add the Vendor ID and Product ID to the "device-type-scsi-options-list" parameter. See the instructions in the driver configuration file and the excerpts below:

  Change to "no" as shown:

  **mpxio-disable="no";**

  Add or comment out these lines:

  **device-type-scsi-options-list =**

  **"HP OPEN", "symmetric-option";**

  **symmetric-option = 0x1000000;**

  (There must be exactly 6 spaces between HP and OPEN.)
- Reboot the host after editing the configuration file.
- After you have rebooted and the LDEV has been defined as a LUN to the host, use the **cfgadm** command to display and configure the controller instances for SAN connectivity. (Controller instance (c#) may differ between systems, and you need to replace the WWPN in the example below with the WWPNs for your array ports):

```
# cfgadm -al
Ap_Id                Type       Receptacle Occupant    Condition
c3                   fc-fabric  connected  configured  unknown
c3::50060e8003285301 disk       connected  configured  unknown
c4                   fc-fabric  connected  configured  unknown
c4::50060e8003285311 disk       connected  configured  unknown
```

**# cfgadm -c configure c3::50060e8003285301**

**# cfgadm -c configure c4::50060e8003285311**

## Configuring JNI HBAs

Configure 2 Gbit and 1 Gbit HBAs differently as explained below. JNI HBA drivers are available from AMCC/JNI at http://www.jni.com/drivers. The JNI EZFibre HBA configuration utility does not allow dynamic LUN addition. Instead edit the configuration files manually as explained below, and use EZFibre for view only purposes. If you do not plan to use dynamic LUN addition, you may use EZFibre to configure the HBA.

Persistent bindings are necessary in a fabric topology and are used to bind a SCSI target ID to a particular WWPN (of an array port). This is required to guarantee that the SCSI target IDs will remain the same when the system is rebooted. Persistent bindings can be set by editing the configuration file as shown in the examples that follow. Make sure the target in the driver configuration file and in the kernel file (/kernel/drv/sd.conf) match. Replace the WWNs shown in the examples with the correct WWNs for your array ports. You can view port WWNs using Command View or LUN Configuration Manager.

## JNI 2 Gbit HBAs

- Use JNIC146x driver (version 5.3 or later).
- Edit the HBA driver settings in the **/kernel/drv/jnic146x.conf** file.

  ### For a SAN environment:

  **FcLoopEnabled = 0;**
  **FcFabricEnabled = 1;**

  ### For direct connect:

  **FcLoopEnabled = 1;**
  **FcFabricEnabled = 0;**

  ### Enable failover for VxVM Dynamic Multi Pathing (DMP):

  **FailoverDelay = 30;**

  ### SAN persistent binding:

  **automap = 0;**
  **jnic146x0-target20_hba = "jnic146x0";**
  **jnic146x0-target20_wwpn = "50060e8003285301";**
  **jnic146x1-target30_hba = "jnic146x1";**
  **jnic146x1-target30_wwpn = "50060e8003285311";**

  **(**Replace the WWPNs above with the WWPNs for your array ports.)
- Make sure the jnic146x entry below is present at the bottom of the **/kernel/drv/sd.conf** file:

  **name="sd" parent="jnic146x" target=0;**
- Perform a reconfiguration reboot so the host can implement the changes you made in the configuration files.
- After configuring LUNS as explained in “Defining the paths” on page 33, use the **jnic146x_update_drv** command to request that the jnic146x driver instances perform LUN rediscovery:

  **`/opt/JNIC146x/jnic146x_update_drv -r -a`**

### JNI 1 Gbit HBAs (JNIC driver, version 4.x)

- Edit the **/kernel/drv/jnic.conf** file:

  ### Set the Fibre topology for a SAN or direct connect:
  ### For a SAN environment:
  **FcLoopEnabled=0;**
  **FcFabricEnabled=1;**
  **FcPortCfgEnable = 1;**
  ### For direct connect:
  **FcLoopEnabled=1;**
  **FcFabricEnabled=0;**
  **FcPortCfgEnable = 0;**

  ### Enable failover for VxVM Dynamic Multi Pathing (DMP):
  **FailoverDelay=30;**

  ### SAN persistent binding:
  **def_hba_binding="null";**
  **target20_wwpn="50060e8003285301";**
  **target20_hba="jnic0";**
  **target30_wwpn="50060e8003285311";**
  **target30_hba="jnic1";**

  **(**Replace the WWPNs above with the WWPNs for your array ports.)

- Edit the **/kernel/drv/sd.conf** file:

  **name="sd" class="scsi" target=20 lun=1;**
  **name="sd" class="scsi" target=20 lun=2;**
  **...**
  **name="sd" class="scsi" target=30 lun=1;**
  **name="sd" class="scsi" target=30 lun=2;**
  **...**

- Perform a reconfiguration reboot to implement the changes to the configuration files.
- If LUNs have been preconfigured in the /kernel/drv/sd.conf file, use the **devfsadm** command to perform LUN rediscovery after configuring LUNS as explained in "Defining the paths" on page 33.

### JNI 1 Gbit HBAs (JNIfcaw/JNIfcaPCI drivers, version 2.x)

- Edit the **/kernel/drv/fca-pci.conf** (or fcaw.conf) file:

  ### Set the Fibre topology for SAN or direct connect:
  ### For a SAN:
  **fca_nport = 1;**
  ### For direct connect:
  **fca_nport = 0;**

  ### Enable failover for VxVM Dynamic Multi Pathing (DMP):
  **failover = 30;**

  ### SAN persistent binding:
  **def_hba_binding = "null";**
  **target20_wwpn="50060e8003285301";**
  **target20_hba="fca-pci0";**
  **target30_wwpn="50060e8003285311";**
  **target30_hba="fca-pci1";**

  **(**Replace the WWPNs above with the WWPNs for your array ports.)

- Edit the **/kernel/drv/sd.conf** file:

  **name="sd" class="scsi" target=20 lun=1;**
  **name="sd" class="scsi" target=20 lun=2;**
  **...**
  **name="sd" class="scsi" target=30 lun=1;**
  **name="sd" class="scsi" target=30 lun=2;**
  **...**

- Perform a reconfiguration reboot to implement the changes to the configuration files.
- If LUNs have been preconfigured in the /kernel/drv/sd.conf file, use the **devfsadm** command to perform LUN rediscovery after configuring LUNS as explained in "Defining the paths" on page 33.

## Configuring Emulex HBAs

Configure Emulex HBA as follows:

- Edit the **/kernel/drv/lpfc.conf** driver configuration file to set up the HBA for a SAN infrastructure:

  **topology = 2;**
  **scan-down = 0;**

- If multiple HBAs and VxVM are used, adjust the following parameters to assure correct VxVM behavior:

  **no-device-delay=0;**
  **nodev-tmo=30;**
  **linkdown-tmo=30;# verify, should be default value**

- Persistent bindings are necessary in a fabric topology and are used to bind a SCSI target ID to a particular WWPN (of an array port). This is required to guarantee that the SCSI target IDs will remain the same when the system is rebooted. Persistent bindings can be set by editing the configuration file or by using the **lputil** utility. The following example illustrates the binding of target 20 (lpfc instance 2) to WWPN 50060e8003285301 and the binding of target 30 (lpfc instance 0) to WWPN 50060e8003285311:

  **fcp-bind-WWPN="50060e8003285301:lpfc2t20",**

  **"50060e8003285311:lpfc0t30";**

  **(**Replace the WWPNs above with the WWPNs for your array ports.)

- For each LUN that needs to be accessed, add an entry to the **/kernel/drv/sd.conf** file. For example, assume you want to access

LUNs 1 and 2 through both paths. You would add the following entries (preferably at the end of the file):

**name="sd" parent="lpfc" target=20 lun=1;**
**name="sd" parent="lpfc" target=20 lun=2;**
**name="sd" parent="lpfc" target=30 lun=1;**
**name="sd" parent="lpfc" target=30 lun=2;**

- Perform a reconfiguration reboot to implement the changes to the configuration files.
- If LUNs have been preconfigured in the /kernel/drv/sd.conf file, use the **devfsadm** command to perform LUN rediscovery after configuring LUNS as explained in "Defining the paths" on page 33.

## Verifying the HBA configuration

After installing the HBAs, verify recognition of the HBAs and drivers.

1. Log into the system as **root**. Verify that all devices are powered on and properly connected to the system.
2. Use the **prtdiag** command (/usr/platform/sun4u/sbin/prtdiag) to verify that the HBA is installed properly. By using the **prtconf** command and/or by browsing the /var/ad/messages file, you can check whether the HBA driver has attached. Look for the WWN/WWPN of the HBA in the /var/adm/messages file or by using an HBA-specific tool or command.

## Clustering and Fabric zoning

If you plan to use clustering, install and configure the clustering software on the servers. Clustering is the organization of multiple servers into groups. Within a cluster, each server is a node. Multiple clusters compose a multi-cluster environment. The following example shows a multi-cluster environment with three clusters, each containing two nodes. The nodes share access to the disk array.

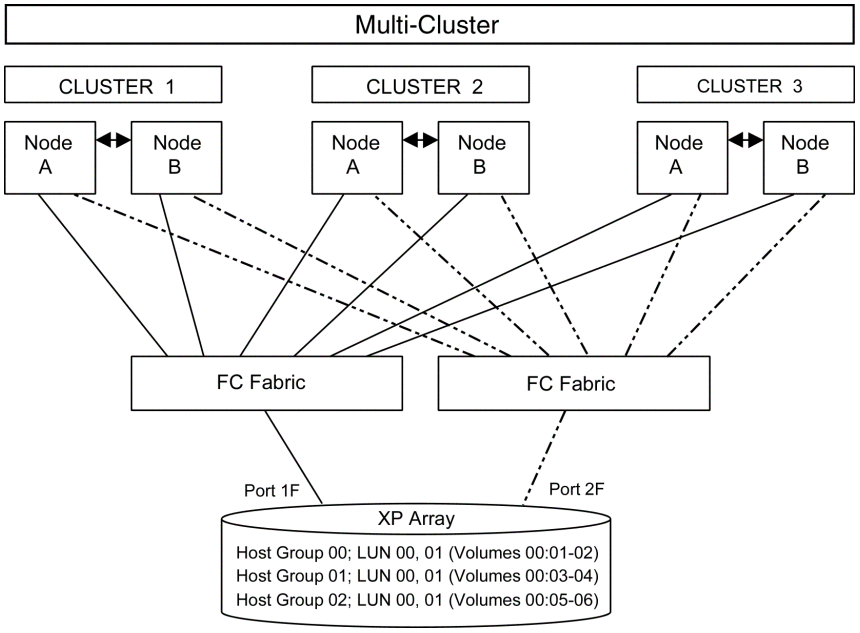


Within the Storage Area Network (SAN), the clusters may be homogeneous (all the same operating system) or they may be heterogeneous (mixed operating systems). How you configure LUN Security and fabric zoning depends on the operating system mix and the SAN configuration.

## Fabric zoning and LUN security for multiple operating systems

By using appropriate zoning and LUN security, you can connect multiple clusters of various operating systems to the same switch and fabric:

- Host zones must contain only homogeneous operating systems.
- Storage port zones may overlap if more than one operating system needs to share an array port.
- Heterogeneous operating systems may share an XP array port if you use Secure Manager and set the appropriate host group and mode; all others must connect to a dedicated XP array port.
- Use Secure Manager for LUN isolation when multiple hosts connect through a shared array port. Secure Manager provides LUN security by allowing you to restrict which LUNs each host can access.

| Environment | OS Mix | Fabric Zoning | LUN Security |
|---|---|---|---|
| Standalone SAN (non-clustered) | homogeneous (a single OS type present in the SAN) | Not required | Must be used when multiple hosts connect through a shared port |
| | heterogeneous (more than one OS type present in the SAN) | Required | |
| Clustered SAN | homogeneous (a single OS type present in the SAN) | Not required | Must be used when multiple cluster nodes connect through a shared port |
| | heterogeneous (more than one OS type present in the SAN) | Required | |
| Multi-Cluster SAN | homogeneous (a single OS type present in the SAN) | Not required | Must be used when multiple cluster nodes connect through a shared port |
| | heterogeneous (more than one OS type present in the SAN) | Required | |

# Connect the disk array

Connect the disk array to the host as follows:

1. The HP service representative verifies operational status of the disk array channel adapters, LDEVs, and paths.
2. The HP representative connects the Fibre Channel cables between the disk array and the host.
3. Verify the ready status of the disk array and peripherals.

# Defining the paths

Use Command View (shown) or LUN Configuration Manager to create paths (LUNs) between hosts and volumes in the disk array, also called LUN mapping. LUN mapping includes these tasks:

- Configuring ports
- Setting LUN security
- Creating host groups by operating system and setting their host modes
- Assigning host bus adapter WWNs to host groups.
- Mapping volumes to host groups (by assigning LUNs).

For details, see the Command View or LUN Configuration Manager guide. HP recommends that you note LUNS and their ports, WWNs, nicknames, and LDEVs for later use in verifying host and device configuration.

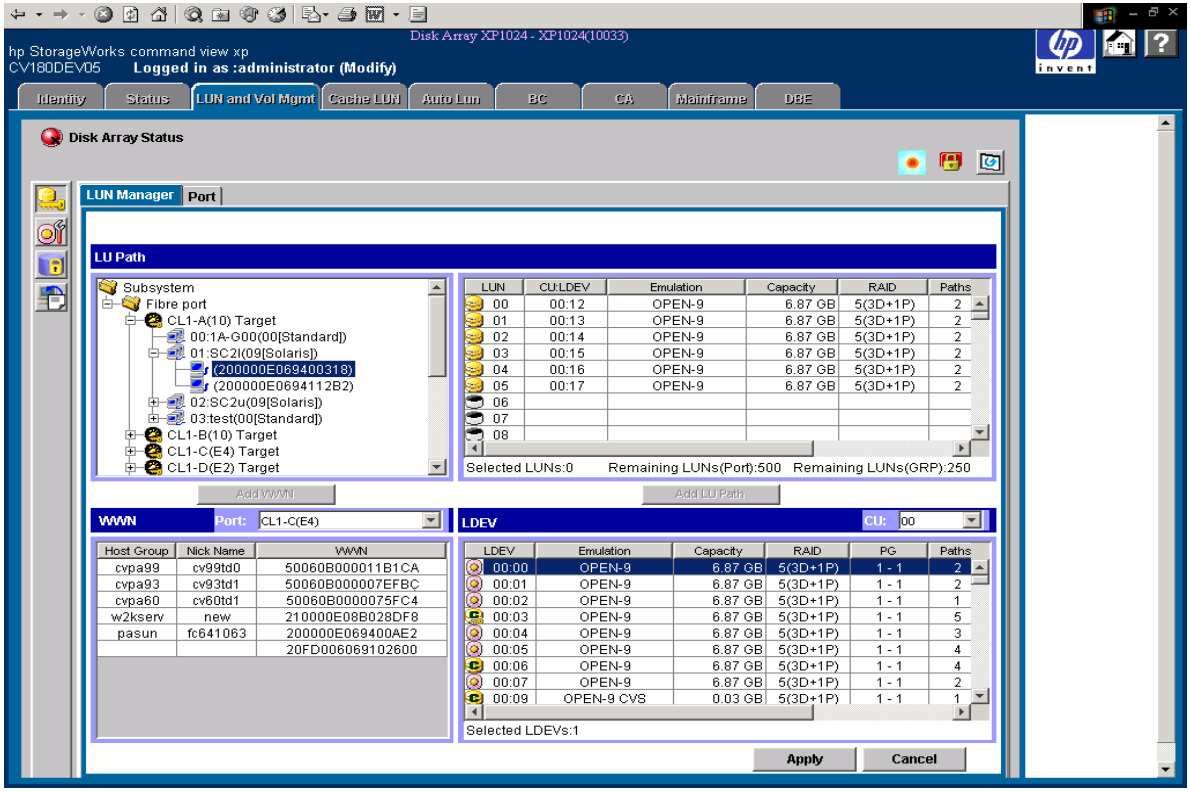

## Adding the new device paths to the system

After "Configuring the HBAs" (page 21), some HBA drivers require you to add the new paths to the SCSI disk driver configuration file on the host (**/kernel/drv/sd.conf**). This enables the Solaris system to recognize the new devices on bootup after the disk array is connected.

**Caution** *To ensure that the system can boot properly even if you make a mistake in the driver configuration file, add the new paths at the end of the file. (This ensures the system boot entries higher up in the file execute first.)*

To avoid reboots when you add new LUNs later, preconfigure extra LUNs in the /kernel/drv/sd.conf file. However, use LUNs sequentially without skipping any numbers to prevent delays during bootup.

After you modify the /kernel/drv/sd.conf file, do a reconfiguration reboot (touch /reconfigure; shutdown -6).

If you added unused LUNs for future use, when you later add those LUNs to the array, use the **devfsadm** command on the host to enable host recognition of the new LUNs.

## Verifying disk array device recognition

Verify that the host server recognizes the disk array devices:

1. Use **format** to display the device information.
2. Check the list of disks to verify the host recognizes all disk array devices. If any devices are missing or if no array devices are shown, check the following:
   - SAN (zoning configuration)
   - Disk array path configuration (HBA WWNs)
   - Host HBA configuration (WWN information, driver instance, target and LUN assignment, and /var/adm/messages)
   - Check cables

# Configure disk array devices

Configure the disk array devices using the procedures that follow. The process is similar but not identical to the steps you would follow to configure any new disk. Creating scripts to configure all devices at once may save you considerable time.

**Configuring devices typically requires these steps:**

1. Initialize, partition, and label each device by using the disk management utility provided with the host server.
2. Create a file system for each device using the **newfs** or **mkfs -F vxfs** command.
3. Create a mount directory for each device using the **mkdir** command.
4. Enter each device into the mount table by editing **/etc/vfstab**.
5. Use the **mount -a** command to auto-mount devices.
6. Use a **df -k** command to verify all devices auto-mounted.

## Partitioning and labeling the devices

Partition and label the new devices using the Sun **format** utility. OPEN-*x* devices can have more than one partition.

**Caution** *Be careful when using the Sun format utility. Some commands are incompatible with the disk array and can overwrite data.*

To partition and label the disks, work with all devices of one type (for example, OPEN-3), then all devices of the next type (for example, OPEN-9), and so on until you have partitioned and labeled all new devices.

During disk partitioning and labeling, enter the disk parameters from the device specification table in Appendix A.

1. Enter **format** at the root prompt to start the format utility.
2. Verify that all new devices are displayed. If they are not, exit the format utility (**quit** or **Ctrl-D**), and make sure the Fibre Channel-to-LDEV paths were defined correctly for all devices and that all new devices were added to the driver configuration file.
3. Record the character-type device file names (for example, c1t2d0) for all of the new disks. You will use this data to create the file systems.
4. When you are asked to specify the disk, enter the number of the device to be partitioned and labeled.
5. When you are asked if you want to label the disk, enter **n** for "no."
6. After the format menu is displayed, enter **type** to display the disk types.
7. If the disk type for the selected device is already defined, enter the number for that disk type.
8. If the disk type for the selected device is *not* already defined, enter the number for "other" to define the disk type.
9. Enter the disk type parameters for the selected device using the Device Geometry tables in Appendix A. See the following example.

```
# format> type
ç Enter type.
:
AVAILABLE DRIVE TYPES
        0. Auto configure
             :
```

```
        14. SUN2.1G
        15. HITACHI-OPEN-3-0315          Do not select this disk type.
        16. other
Specify disk type (enter its number):16  Enter number for "other"
Enter number of data cylinders:3336      Enter value from Appendix A
Enter number of alternate cylinders[2]:2 Enter value from Appendix A
Enter number of physical cylinders[3338]:Press Enter for default
Enter number of heads:15                 Enter value from Appendix A
Enter number of physical sectors/track[defaults]: Press Enter for default
Enter rpm of drive [3600]:10000          Enter value from Appendix A
Enter format time[defaults]:             Press Enter for default
Enter cylinder skew[defaults]:           Press Enter for default
Enter track skew[defaults]:              Press Enter for default
Enter track per zone[defaults]:          Press Enter for default
Enter alternate tracks[defaults]:        Press Enter for default
Enter cache control[defaults]:           Press Enter for default
Enter prefetch threshold[defaults]:      Press Enter for default
Enter minimum prefetch[defaults]:        Press Enter for default
Enter maximum prefetch[defaults]:        Press Enter for default
Enter disk type name:"HITACHI OPEN-3"    Enter device type from Appendix A
selecting c1t2d0
[disk formatted]
No defined partition tables.
Disk not labeled.  Label it now ? n      Enter "n" for no.
format>
```

10. If you **do not** want to partition and label the disk, skip to step 16. (You don't need to partition or label if you use Veritas Volume Manager.)

    If you **do** want to partition the disk, when you are asked whether you want to label the disk, enter **n** for "no."

11. After the format menu is displayed, enter **partition**.
12. Enter the desired partition number and partition parameters.
13. Enter **print** at the partition> prompt to display the partition table.
14. When you are finished setting the partitions for the selected device, enter **label** at the partition> prompt, and then enter **y** to label the device.
15. Exit the partition utility and return to the format utility.
16. Display the available disks by entering **disk** at the **format>** prompt. Make sure the disk you just labeled is correctly displayed.
17. Repeat this labeling procedure for each new device.
18. When you finish labeling the disks and verifying the disk labels, enter **quit** or press **Ctrl-D** to exit the format utility.

## Creating the file systems

1. If you want to create a UFS file system, create the file system using the **newfs –C** *maxcontig* command.

*Example*

```
# newfs -C 6 /dev/rdsk/c1t2d0s0
```

Use 6 or one of the following multiples of 6 as the *maxcontig* value for all disk array OPEN-*x* devices: 12, 18, 24, or 30.

If you use 6, the Sun OS will access 48 KB as a unit (6*8 KB), which matches the track size of the OPEN-*x* devices. These *maxcontig* values (6, 12, 18, 24, 30) optimize the I/O performance of the disk array by keeping the I/O data range on one track. The *maxcontig* value that you choose depends on your applications, and you can change the *maxcontig* parameter to a different value at any time.

Use the character-type device file (for example, **/dev/rdsk/c1t2d0s0**) as the argument.

2. When the confirmation appears, enter **y** for yes if the file name is correct. If the file name is not correct, enter **n** and repeat step 1.
3. Repeat this procedure for each new OPEN-*x* device. Use the same *maxcontig* value for all disk array devices.
4. Create mount directories and set up auto-mounting for all disks as you normally would.

# Configure for use with Veritas Volume Manager

HP XP disk arrays are certified for VxVM support.

Be sure the failover parameter is set correctly when you install the HBA. Failure to do so results in a loss of path failover in Dynamic Multipathing (DMP). See "Configuring the HBAs" on page 21 and the HBA manufacturer's instructions for specific HBA parameters to set.

## VxVM versions prior to 3.2

To ensure proper operation of the DMP feature of VxVM, update the **/kernel/drv/vxdmp.conf** file to reflect the appropriate Vendor string.

*Example*

```
name="vxdmp" parent="pseudo" instance=0
dmp_jbod="HP";
```

## VxVM versions 3.2 and later

More recent versions of VxVM use Array Support Libraries (ASL) to set up the Dynamic Multipathing (DMP) feature and other parameters. The ASL is required only for XP128 and XP1024 arrays.

You can download the ASL by searching for "XP128" on the Veritas website: http://support.veritas.com. ASL installation instructions are supplied with the Volume Manager and are also available on the Veritas website.

# 2

# Troubleshooting

If you encounter an error condition, see “Error conditions” on page 42 for recommended actions.

If you are unable to resolve an error condition, ask your HP support representative for assistance. See “Calling the HP support center” on page 44.

# Error conditions

| Error Condition | Recommended Action |
|---|---|
| The logical devices are not recognized by the host. | Verify that the READY indicator lights on the disk array are ON.<br><br>Verify that fiber cables are correctly installed and firmly connected.<br><br>Verify that the target IDs are properly configured and either match the loop ID in a direct environment or match the target assignment from the driver configuration file in a SAN environment. For drivers where sd.conf needs to be configured, the LUNs for each target should start at 1 and continue sequentially without skipping numbers. Verify that /kernel/drv/sd.conf is correctly configured. For drivers where sd.conf is not configured, the LUNs for each target should start at 0 and continue sequentially without skipping numbers.<br><br>Check HBA configuration as explained in Chapter 1. Look for the HBA WWN in the /var/adm/messages file and verify that it matches the WWN configured on the XP array. Verify the correct WWN for the XP ports have been configured in the driver configuration file (WWN binding). Verify the correct driver instance has been configured for the given WWN.<br><br>Use the **prtdiag** command to verify the HBA is properly installed (/usr/platform/sun4u/sbin/prtdiag). With the prtconf command and/or by browsing the /var/adm/messages file, you can check that the HBA driver has attached.<br><br>Check SAN (zoning configuration), which is explained in Chapter 1.<br><br>A reconfiguration reboot or a rescan for new LUNs may be required (**devfsadm** command or for JNI HBAs **jnic146x_update_drv -r -a**).<br><br>Verify that the disk array Host Mode is set correctly. |

| Error Condition | Recommended Action |
|---|---|
| The host does not reboot properly after hard shutdown. | If you power off the host without executing the shutdown process, wait three minutes to allow the disk array's internal timeout process to purge queued commands. If the host restarts while the disk array is processing queued commands, the host may not reboot successfully. |
| Physical volumes cannot be created. | Verify that the disk array logical devices are correctly formatted. |
| Disks are not visible in VxVM | Verify that the LUN has been correctly labeled by the **format** command. |
| A file system cannot be created. | Check that a character-type device file is specified as the device file in the **newfs** command.<br><br>Verify that the logical unit is correctly labeled by the **format** command. |
| A file system is not mounted after rebooting. | Verify that the host was restarted correctly.<br><br>Verify that the file system attributes are correct.<br><br>Verify that **/etc/vfstab** was edited correctly. |
| The disk array performs a self reboot because the disk array was busy or it logged a panic message. | Reboot the host. |
| The disk array responds "Not Ready" or the disk array has displayed "Not Ready" and timed out. | Contact HP. |
| The host detects a parity error. | Check the HBA and make sure it was installed properly.<br><br>Reboot the host. |
| HBA driver is not attached. | Verify the installed HBA driver is the correct one for the HBA and the HBA firmware. |
| vxinstall aborts after device detection. | Verify that you have installed the Array Support Library. |

# Calling the HP support center

If you need to call HP customer support, provide as much information about the problem as possible, including the circumstances of the error or failure and the exact content of any error messages.

Depending on your system configuration, you may be able to view error messages as follows:

- View SIMs in Command View (Device Health tab).
- View R-SIMs in Remote Control XP, including reference codes and severity levels of recent R-SIMs.
- View SIMs that generate SNMP traps on the host.

# A

# Disk array device emulations

This appendix provides information about disk array supported emulations and device type specifications. Some parameters may not be relevant to your array. Consult your HP representative for information about supported configurations for your system.

# Supported emulations

| XP Type | Emulation | OPEN-x | LUSE | CVS | LUSE & CVS |
|---|---|---|---|---|---|
| XP48<br>XP512 | OPEN-3 | Yes | Yes | Yes | Yes |
| | OPEN-8 | Yes | Yes | Yes | Yes |
| | OPEN-9 | Yes | Yes | Yes | Yes |
| | OPEN-E | Yes | Yes | Yes | Yes |
| | OPEN-K | Yes | Yes | Yes | Yes |
| | OPEN-L | Yes | Yes | | |
| | OPEN-M | Yes | Yes | | |
| | OPEN-V | | | | |
| XP128<br>XP1024<br>XP12000 | OPEN-3 | Yes | Yes | Yes | Yes |
| | OPEN-8 | Yes | Yes | Yes | Yes |
| | OPEN-9 | Yes | Yes | Yes | Yes |
| | OPEN-E | Yes | Yes | Yes | Yes |
| | OPEN-K | | | | |
| | OPEN-L | Yes | Yes | | |
| | OPEN-M | | | | |
| | OPEN-V | Yes | Yes | | |

# Device type specifications

| **Device Type** *(Note 1)* | **Category** *(Note 2)* | **Blocks (512 bytes )** | **Sector Size (bytes)** | **# of Cylinders** | **Heads** | **Sectors per Track** | **Capacity MB*** *(Note 3)* |
|---|---|---|---|---|---|---|---|
| OPEN-3 | SCSI disk | 4806720 | 512 | 3338 | 15 | 96 | 2347 |
| OPEN-8 | SCSI disk | 14351040 | 512 | 9966 | 15 | 96 | 7007 |
| OPEN-9 | SCSI disk | 14423040 | 512 | 10016 | 15 | 96 | 7042 |
| OPEN-E | SCSI disk | 28452960 | 512 | 19759 | 15 | 96 | 13893 |
| OPEN-L | SCSI disk | 71192160 | 512 | 49439 | 15 | 96 | 34761 |
| OPEN-V | SCSI disk | max=125827200 | 512 | Note 5 | 15 | 128 | Note 6 |
| | | | | | | | **LUSE** |
| OPEN-3*n | SCSI disk | 4806720*n | 512 | 3338*n | 15 | 96 | 2347*n |
| OPEN-8*n | SCSI disk | 14351040*n | 512 | 9966*n | 15 | 96 | 7007*n |
| OPEN-9*n | SCSI disk | 14423040*n | 512 | 10016*n | 15 | 96 | 7042*n |
| OPEN-E*n | SCSI disk | 28452960*n | 512 | 19759*n | 15 | 96 | 13893*n |
| OPEN-L*n | SCSI disk | 71192160*n | 512 | 49439*n | 15 | 96 | 34761*n |
| OPEN-V*n | SCSI disk | max=125827200 Note 4 | 512 | Note 5 | 15 | 128 | Note 6 |
| | | | | | | | **CVS** |
| OPEN-3 CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-8 CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-9 CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-E CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| | | | | | | | **CVS LUSE** |
| OPEN-3*n CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-8*n CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-9*n CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-E*n CVS | SCSI disk | Note 4 | 512 | Note 5 | 15 | 96 | Note 6 |
| OPEN-V*n | SCSI disk | Note 4 | 512 | Note 5 | 15 | 128 | Note 6 |

*Capacity = (512 x number of blocks) ÷ $1024^2$

***Note 1:*** The availability of a disk type depends on the disk array.

***Note 2:*** The devices are defined to the host as SCSI disk devices, even though the interface is Fibre Channel.

***Note 3*** The device capacity can sometimes be changed by the BIOS or host adapter board. This may make actual capacity different from that listed in the table.

***Note 4:*** The number of blocks for a CVS volume is calculated as follows:
# of blocks = (# of cylinders) × (# of heads) × (# of sectors per track)

***Example 1:*** For an OPEN-3 CVS volume with capacity = 37 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (96 sectors per track) = 76320

***Example 2:*** For an OPEN-V CVS volume with capacity = 49 MB:
# of blocks = (53 cylinders–see Note 5) × (15 heads) × (128 sectors per track) = 101760

***Note 5:*** The number of cylinders for a CVS volume is calculated as follows (↑ …↑ means that the value should be rounded up to the next integer):

**OPEN-3/8/9/E:** The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑

***Example:*** For an OPEN-3 CVS volume with capacity = 37 MB:
# of cylinders = ↑ 37 × 1024/720 ↑ = ↑ 52.62 ↑ (rounded up to next integer) = 53 cylinders

**OPEN-V:** The number of cylinders for a CVS volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑

***Example:*** For an OPEN-V CVS volume with capacity = 49 MB:
# of cylinders = ↑ 49 × 16/15 ↑ = ↑ 52.26 ↑ (rounded up to next integer) = 53 cylinders

**OPEN-3/8/9/E:** The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 1024/720 ↑ × n

***Example:*** For a CVS LUSE volume with capacity = 37 MB and n = 4
# of cylinders = ↑ 37 × 1024/720↑ × 4 = ↑ 52.62↑ × 4 = 53 × 4 = 212

**OPEN-V:** The number of cylinders for a CVS LUSE volume =
# of cylinders = ↑ (capacity (MB) specified by user) × 16/15 ↑ × n

***Example:*** For an OPEN-V CVS LUSE volume with capacity = 49 MB and n = 4
# of cylinders = ↑ 49 × 16/15↑ × 4 = ↑ 52.26↑ × 4 = 53 × 4 = 212

***Note 6:*** The capacity of an OPEN-3/8/9/E CVS volume is specified in MB, not number of cylinders. The capacity of an OPEN-V CVS volume can be specified in MB or number of cylinders. You set the volume size using the LUN Configuration Manager or Command View software.

# Device Geometry Tables

Use the following geometry tables when configuring devices and file systems on the host.

## OPEN-x geometry table

| Device Type | # of Data Cylinders | # of Alternate Cylinders | RPM | Partition Size (sample) |
|---|---|---|---|---|
| OPEN-3 | 3336 | 2 | 10,000 | 3336c |
| OPEN-8 | 9964 | 2 | 10,000 | 9964c |
| OPEN-9 | 10014 | 2 | 10,000 | 10014c |
| OPEN-E | 19757 | 2 | 10,000 | 19757c |
| OPEN-L | 19013 | 2 | 10,000 | 19013c |
| OPEN-3*n | N1* | 2 | 10,000 | N4* |
| OPEN-8*n | N26* | 2 | 10,000 | N29* |
| OPEN-9*n | N5* | 2 | 10,000 | N8* |
| OPEN-E*n | N30* | 2 | 10,000 | N33* |
| OPEN-L*n | N34 | 2 | 10,000 | N37 |
| OPEN-*x* VLL | See Device Type Specs | 2 | 10,000 | See Device Type Specs |
| OPEN-3*n VLL | N22* | 2 | 10,000 | N25* |
| OPEN-8*n VLL | N22* | 2 | 10,000 | N25* |
| OPEN-9*n VLL | N22* | 2 | 10,000 | N25* |
| OPEN-E*n VLL | N22* | 2 | 10,000 | N25* |
| OPEN-V*n VLL | N22* | 2 | 10,000 | N25* |

* For the values indicated by Nxx (for example, N15, N22), refer to the tables that follow.

## OPEN-3*n LUSE device geometry table

| n | **Data Cylinders** N1 **Partition Size** N4 | **Heads** N2 | **Blocks/Track** N3 | **Usable Blocks** (N1+2)*N2*N3 | **Provided Blocks** 3338*15*96*n | **Diff.** |
|---|---|---|---|---|---|---|
| 2 | 6674 | 15 | 96 | 9613440 | 9613440 | 0 |
| 3 | 10012 | 15 | 96 | 14420160 | 14420160 | 0 |
| 4 | 13350 | 15 | 96 | 19226880 | 19226880 | 0 |
| 5 | 16688 | 15 | 96 | 24033600 | 24033600 | 0 |
| 6 | 20026 | 15 | 96 | 28840320 | 28840320 | 0 |
| 7 | 23364 | 15 | 96 | 33647040 | 33647040 | 0 |
| 8 | 26702 | 15 | 96 | 38453760 | 38453760 | 0 |
| 9 | 30040 | 15 | 96 | 43260480 | 43260480 | 0 |
| 10 | 16688 | 30 | 96 | 48067200 | 48067200 | 0 |
| 11 | 20026 | 33 | 80 | 52873920 | 52873920 | 0 |
| 12 | 20026 | 30 | 96 | 57680640 | 57680640 | 0 |
| 13 | 20026 | 39 | 80 | 62487360 | 62487360 | 0 |
| 14 | 23364 | 30 | 96 | 67294080 | 67294080 | 0 |
| 15 | 16688 | 45 | 96 | 72100800 | 72100800 | 0 |
| 16 | 26702 | 30 | 96 | 76907520 | 76907520 | 0 |
| 17 | 30040 | 34 | 80 | 81714240 | 81714240 | 0 |
| 18 | 30040 | 30 | 96 | 86520960 | 86520960 | 0 |
| 19 | 30040 | 38 | 80 | 91327680 | 91327680 | 0 |
| 20 | 16688 | 60 | 96 | 96134400 | 96134400 | 0 |
| 21 | 23364 | 45 | 96 | 100941120 | 100941120 | 0 |
| 22 | 30040 | 55 | 64 | 105747840 | 105747840 | 0 |
| 23 | 30040 | 46 | 80 | 110554560 | 110554560 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 24 | 20026 | 60 | 96 | 115361280 | 115361280 | 0 |
| 25 | 16688 | 45 | 160 | 120168000 | 120168000 | 0 |
| 26 | 20026 | 39 | 160 | 124974720 | 124974720 | 0 |
| 27 | 30040 | 45 | 96 | 129781440 | 129781440 | 0 |
| 28 | 23364 | 60 | 96 | 134588160 | 134588160 | 0 |
| 29 | 30040 | 58 | 80 | 139394880 | 139394880 | 0 |
| 30 | 16688 | 45 | 192 | 144201600 | 144201600 | 0 |
| 31 | 30040 | 62 | 80 | 149008320 | 149008320 | 0 |
| 32 | 26702 | 60 | 96 | 153815040 | 153815040 | 0 |
| 33 | 30040 | 55 | 96 | 158621760 | 158621760 | 0 |
| 34 | 30040 | 64 | 85 | 163428480 | 163428480 | 0 |
| 35 | 30040 | 56 | 100 | 168235200 | 168235200 | 0 |
| 36 | 30040 | 60 | 96 | 173041920 | 173041920 | 0 |

N1, N2, N3: Use values in OPEN-x specification tables on pages 47 and 50.

N4: Use same value as N1. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 6674c for OPEN-3*2).

## OPEN-8*n LUSE device geometry table

| n | **Data Cylinders** N26 **Partition Size** N29 | **Heads** N27 | **Blocks/Track** N28 | **Usable Blocks** (N26+2)*N27*N28 | **Provided Blocks** 9966*15*96*n | **Diff.** |
|---|---|---|---|---|---|---|
| 2 | 19930 | 15 | 96 | 28702080 | 28702080 | 0 |
| 3 | 29896 | 15 | 96 | 43053120 | 43053120 | 0 |
| 4 | 29896 | 20 | 96 | 57404160 | 57404160 | 0 |
| 5 | 29896 | 25 | 96 | 71755200 | 71755200 | 0 |
| 6 | 29896 | 30 | 96 | 86106240 | 86106240 | 0 |
| 7 | 29896 | 35 | 96 | 100457280 | 100457280 | 0 |
| 8 | 29896 | 40 | 96 | 114808320 | 114808320 | 0 |
| 9 | 29896 | 45 | 96 | 129159360 | 129159360 | 0 |
| 10 | 29896 | 50 | 96 | 143510400 | 143510400 | 0 |
| 11 | 29896 | 55 | 96 | 157861440 | 157861440 | 0 |
| 12 | 29896 | 60 | 96 | 172212480 | 172212480 | 0 |
| 13 | 29896 | 52 | 120 | 186563520 | 186563520 | 0 |
| 14 | 29896 | 56 | 120 | 200914560 | 200914560 | 0 |
| 15 | 29896 | 60 | 120 | 215265600 | 215265600 | 0 |
| 16 | 29896 | 64 | 120 | 229616640 | 229616640 | 0 |
| 17 | 29896 | 34 | 240 | 243967680 | 243967680 | 0 |
| 18 | 29896 | 36 | 240 | 258318720 | 258318720 | 0 |
| 19 | 29896 | 38 | 240 | 272669760 | 272669760 | 0 |
| 20 | 29896 | 40 | 240 | 287020800 | 287020800 | 0 |
| 21 | 29896 | 42 | 240 | 301371840 | 301371840 | 0 |
| 22 | 29896 | 44 | 240 | 315722880 | 315722880 | 0 |
| 23 | 29896 | 46 | 240 | 330073920 | 330073920 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 24 | 29896 | 48 | 240 | 344424960 | 344424960 | 0 |
| 25 | 29896 | 50 | 240 | 358776000 | 358776000 | 0 |
| 26 | 29896 | 52 | 240 | 373127040 | 373127040 | 0 |
| 27 | 29896 | 54 | 240 | 387478080 | 387478080 | 0 |
| 28 | 29896 | 56 | 240 | 401829120 | 401829120 | 0 |
| 29 | 29896 | 58 | 240 | 416180160 | 416180160 | 0 |
| 30 | 29896 | 60 | 240 | 430531200 | 430531200 | 0 |
| 31 | 29896 | 62 | 240 | 444882240 | 444882240 | 0 |
| 32 | 29896 | 64 | 240 | 459233280 | 459233280 | 0 |
| 33 | 32614 | 60 | 242 | 473584320 | 473584320 | 0 |
| 34 | 29896 | 64 | 255 | 487935360 | 487935360 | 0 |
| 35 | 30655 | 64 | 256 | 502284288 | 502286400 | 2112 |
| 36 | 31531 | 64 | 256 | 516636672 | 516637440 | 768 |

N26, N27, N28: Use values in OPEN-x specification tables on pages 47 and 50.

N29: Use same value as N26. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 19930c for OPEN-8*2).

Data cylinders must be less than or equal to **32767**, heads must be less than or equal to **64**, blocks per track must be less than or equal to **256** when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-3*2 ~ OPEN-3*36 can be used by above parameters.

# OPEN-9*n LUSE device geometry table

| n | **Data Cylinders** N5 **Partition Size** N8 | **Heads** N6 | **Blocks/Track** N7 | **Usable Blocks** (N5+2)*N6*N7 | **Provided Blocks** 10016*15*96*n | **Diff.** |
|---|---|---|---|---|---|---|
| 2 | 20030 | 15 | 96 | 28846080 | 28846080 | 0 |
| 3 | 30046 | 15 | 96 | 43269120 | 43269120 | 0 |
| 4 | 30046 | 20 | 96 | 57692160 | 57692160 | 0 |
| 5 | 30046 | 25 | 96 | 72115200 | 72115200 | 0 |
| 6 | 30046 | 30 | 96 | 86538240 | 86538240 | 0 |
| 7 | 30046 | 35 | 96 | 100961280 | 100961280 | 0 |
| 8 | 30046 | 40 | 96 | 115384320 | 115384320 | 0 |
| 9 | 30046 | 45 | 96 | 129807360 | 129807360 | 0 |
| 10 | 30046 | 50 | 96 | 144230400 | 144230400 | 0 |
| 11 | 30046 | 55 | 96 | 158653440 | 158653440 | 0 |
| 12 | 30046 | 60 | 96 | 173076480 | 173076480 | 0 |
| 13 | 30046 | 52 | 120 | 187499520 | 187499520 | 0 |
| 14 | 30046 | 56 | 120 | 201922560 | 201922560 | 0 |
| 15 | 30046 | 60 | 120 | 216345600 | 216345600 | 0 |
| 16 | 30046 | 64 | 120 | 230768640 | 230768640 | 0 |
| 17 | 30046 | 34 | 240 | 245191680 | 245191680 | 0 |
| 18 | 30046 | 36 | 240 | 259614720 | 259614720 | 0 |
| 19 | 30046 | 38 | 240 | 274037760 | 274037760 | 0 |
| 20 | 30046 | 40 | 240 | 288460800 | 288460800 | 0 |
| 21 | 30046 | 42 | 240 | 302883840 | 302883840 | 0 |
| 22 | 30046 | 44 | 240 | 317306880 | 317306880 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 23 | 30046 | 46 | 240 | 331729920 | 331729920 | 0 |
| 24 | 30046 | 48 | 240 | 346152960 | 346152960 | 0 |
| 25 | 30046 | 50 | 240 | 360576000 | 360576000 | 0 |
| 26 | 30046 | 52 | 240 | 374999040 | 374999040 | 0 |
| 27 | 30046 | 54 | 240 | 389422080 | 389422080 | 0 |
| 28 | 30046 | 56 | 240 | 403845120 | 403845120 | 0 |
| 29 | 30046 | 58 | 240 | 418268160 | 418268160 | 0 |
| 30 | 30046 | 60 | 240 | 432691200 | 432691200 | 0 |
| 31 | 30046 | 62 | 240 | 447114240 | 447114240 | 0 |
| 32 | 30046 | 64 | 240 | 461537280 | 461537280 | 0 |
| 33 | 30985 | 64 | 240 | 475960320 | 475960320 | 0 |
| 34 | 31924 | 64 | 240 | 490383360 | 490383360 | 0 |
| 35 | 31298 | 63 | 256 | 504806400 | 504806400 | 0 |
| 36 | 31689 | 64 | 256 | 519225344 | 519229440 | 4096 |

N5, N6, N7: Use values in OPEN-x specification tables on pages 47, 50, and 51.

N8: Use same value as N5. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (for example, enter 20030c for OPEN-9*2).

## OPEN-E*n LUSE devices geometry table

| n | **Data Cylinders** N30 **Partition Size** N33 | **Heads** N31 | **Blocks/Track** N32 | **Usable Blocks** (N30+2)*N31*N32 | **Provided Blocks** 9966*15*96*n | **Diff.** |
|---|---|---|---|---|---|---|
| 2 | 19757 | 30 | 96 | 56905920 | 56905920 | 0 |
| 3 | 19757 | 45 | 96 | 85358880 | 85358880 | 0 |
| 4 | 19757 | 60 | 96 | 113811840 | 113811840 | 0 |
| 5 | 19757 | 30 | 240 | 142264800 | 142264800 | 0 |
| 6 | 19757 | 45 | 192 | 170717760 | 170717760 | 0 |
| 7 | 19757 | 60 | 168 | 199170720 | 199170720 | 0 |
| 8 | 19757 | 60 | 192 | 227623680 | 227623680 | 0 |
| 9 | 19757 | 60 | 216 | 256076640 | 256076640 | 0 |
| 10 | 19757 | 60 | 240 | 284529600 | 284529600 | 0 |
| 11 | 27166 | 60 | 192 | 312975360 | 312982560 | 7200 |
| 12 | 29636 | 60 | 192 | 341429760 | 341435520 | 5760 |
| 13 | 32106 | 60 | 192 | 369884160 | 369888480 | 4320 |
| 14 | 27660 | 60 | 240 | 398332800 | 398341440 | 8640 |
| 15 | 29636 | 60 | 240 | 426787200 | 426794400 | 7200 |
| 16 | 31612 | 60 | 240 | 455241600 | 455247360 | 5760 |
| 17 | 31612 | 60 | 255 | 483694200 | 483700320 | 6120 |
| 18 | 31257 | 64 | 256 | 512147456 | 512153280 | 5824 |

N30, N31, N32: Use values in OPEN-x specification tables on pages 47 and 50.

N33: Use same value as N30. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (e.g. enter 19757c for OPEN-E*2).

Data cylinders must be less than or equal to **32767**, heads must be less than or equal to 64, blocks per track must be less than or equal to 256 when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-E*2~OPEN-E*10 can be used by above parameters. About OPEN-E*11~OPEN-E*18, some blocks must become unusable.

## OPEN-L*n LUSE device geometry table

| n | **Data Cylinders** N34 **Partition Size** N37 | **Head** N35 | **Blocks/Track** N36 | **Usable Blocks** (N34+2)*N35*N36 | **Provided Blocks** 49439*15*96*n | **Diff.** |
|---|---|---|---|---|---|---|
| 2 | 19013 | 64 | 117 | 142384320 | 142384320 | 0 |
| 3 | 30422 | 36 | 195 | 213576480 | 213576480 | 0 |
| 4 | 30422 | 45 | 208 | 284768640 | 284768640 | 0 |
| 5 | 30422 | 60 | 195 | 355960800 | 355960800 | 0 |
| 6 | 30422 | 60 | 234 | 427152960 | 427152960 | 0 |
| 7 | 30897 | 63 | 256 | 498339072 | 498345120 | 6048 |

N34, N35, N36: Use values in OPEN-x specification tables on pages 47 and 50.

N37: Use same value as N34. Specify as NNNNc, where NNNN = # of cylinders and c = cylinder (e.g. enter 19013c for OPEN-L*2).

Data cylinders must be less than or equal to **32767**, heads must be less than or equal to 64, blocks per track must be less than or equal to 256 when these values are specified as parameters of Solaris format type subcommand. The whole data blocks of OPEN-L*2~OPEN-L*6 can be used by above parameters. About OPEN-L*7, some blocks must become unusable.

## OPEN-x∗n VLL-LUSE devices geometry table (example)

| # | **Data Cylinders** N22 **Partition Size** N25 | **Head** N23 | **Blocks/Track** N24 | **Usable Blocks** (N22+2)*N23*N24 | **Provided Blocks** N21 | **Diff.** |
|---|---|---|---|---|---|---|
| 1 | 98 | 15 | 96 | 144000 | 35MB×2 volumes<br>↑ 35×1024/720↑ ×2=100<br>100×15×96=144000 | 0 |
| 2 | 2590 | 15 | 96 | 3732480 | 50MB×36 volumes<br>↑ 50×1024/720↑ ×36=2592<br>2592×15×96=3732480 | 0 |
| 3 | 284 | 15 | 96 | 411840 | 100MB×2 volumes<br>↑ 100×1024/720↑ ×2=286<br>286×15×96=411840 | 0 |
| 4 | 5694 | 15 | 96 | 8202240 | 500MB×8 volumes<br>↑ 500×1024/720↑ ×8=5696<br>5696×15×96=8202240 | 0 |
| 5 | 22758 | 30 | 96 | 65548800 | 2000MB×2 volumes<br>↑ 2000×1024/720↑ ×16=45520<br>45520×15×96=65548800 | 0 |
| 6 | 27455 | 40 | 188 | 206476640 | 2800MB×36 volumes<br>↑ 2800×1024/720↑ ×36=143388<br>143388×15×96=206478720 | 2080 |

N21: number of LUSE blocks is calculated as follows:
N21 = N20 × (# of heads) × (# of sectors per track).

N22: N20 – 2, Use total cylinder - 2.

N23, N24: Use values in OPEN-x specification tables on pages 47, 50, and 51.

N25: Use same value as N22.

## OPEN-V*n VLL-LUSE device geometry table (example)

| # | **Data Cylinders** N22 **Partition Size** N25 | **Heads** N23 | **Blocks/Track** N24 | **Usable Blocks** (N22+2)*N23*N24 | **Provided Blocks** N21 | **Diff.** |
|---|---|---|---|---|---|---|
| 1 | 48 | 15 | 128 | 92160 | 45 MB volumes<br>↑ 45 × 16/15 ↑ = 48<br>48 × 15 × 128 = 92160 | 0 |
| 2 | 27305 | 30 | 128 | 104858880 | 50 GB volumes<br>↑ 50 × 1024 × 16/15 ↑ = 54614<br>54614 × 15 × 128 = 104858880 | 0 |
| 3 | 10921 | 150 | 128 | 209721600 | 10 GB × 10 volumes<br>↑ 10× 1024× 16/15↑ × 10=109230<br>109230 × 15 × 128 = 209721600 | 0 |
| 4 | 32767 | 100 | 128 | 419443200 | 20 GB × 10 volumes<br>↑ 20× 1024× 16/15↑ × 10=218460<br>218460 × 15 × 128 = 419443200 | 0 |

N21 number of LUSE blocks is calculated as follows:
N21 = N20 × (# of heads) × (# of sectors per track).

N22: N20 – 2, Use total cylinder - 2.

N23, N24: Use values in OPEN-x specification tables on pages 47, 50, and 51.

N25: Use same value as N22.

# Glossary

| | |
|---|---|
| **AL** | Arbitrated loop. |
| **AL-PA** | Arbitrated loop physical address. |
| **BC** | HP StorageWorks Business Copy XP. BC lets you maintain up to nine local copies of logical volumes on the disk array. |
| **CA** | HP StorageWorks Continuous Access XP. CA lets you create and maintain duplicate copies of local logical volumes on a remote disk array. |
| **Command View** | HP StorageWorks Command View XP, a software product for managing XP arrays. Command View runs on a Windows-based management workstation. |
| **command device** | An LDEV that transfers RAID commands to BC or CA logical volumes. |
| **CVS** | CVS devices (OPEN-x CVS) are custom volumes that are smaller than normal fixed-sized logical disk devices (volumes). |
| **DKC (disk controller unit)** | The array cabinet that houses the channel adapters and service processor (SVP). |
| **DKU (disk cabinet unit)** | The array cabinets that house the disk array physical disks. |
| **emulation modes** | Emulation modes can be assigned to LDEVs to make them operate like standard OPEN system disk drives. The emulation mode of an LDEV determines its capacity. Refer to the appendices for device capacities. |
| **FC** | Fibre Channel. |

| | |
|---|---|
| **FC-AL** | Fibre Channel arbitrated loop. |
| **FCP** | Fibre Channel Protocol. |
| **HBA** | Host bus adapter. |
| **HP** | Hewlett-Packard Company. |
| **LDEV** | Logical device. An LDEV is created when a RAID group is divided into sections using a selected host emulation mode (for example, OPEN-9 or OPEN-M). The number of resulting LDEVs depends on the emulation mode. "LDEV" and "volume" are synonyms. |
| **LUN** | Logical unit number. A LUN results from mapping a SCSI logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN. For example, a LUN associated with two OPEN-3 LDEVs has a size of 4,693 MB. |
| **LUSE** | Logical Unit Size Expansion, a feature which logically combines LDEVs so they appear as a larger LDEV. This allows a LUN to be associated with 2 to 36 LDEVs. LUSE allows applications to access data requiring large amounts of disk space. |
| **OFC** | Open Fibre Control. |
| **OPEN-*x*** | A general term describing any one of the supported OPEN emulation modes (for example, OPEN-L). |
| **OS** | Operating system. |
| **PA** | Physical address. |
| **path** | "Path" and "LUN" are synonymous. Paths are created by associating a port, a target, and a LUN ID with one or more LDEVs. |
| **port** | A connector on a channel adapter card in the disk array. A port passes data between the disk array and external devices, such as a host server. Ports are named using a port group and port letter, for example, CL1-A. |

| | |
|---|---|
| **RAID** | Redundant array of independent disks. |
| **remote console PC** | The PC running HP StorageWorks Remote Control XP. |
| **Remote Control (RC)** | HP StorageWorks Remote Control XP. A software product used for managing XP arrays. |
| **R-SIM** | Remote service information message. |
| **SCSI** | Small computer system interface. |
| **SIM** | Service information message. |
| **SNMP** | Simple Network Management Protocol. |
| **SVP** | Service processor. A notebook computer built into the disk array. The SVP provides a direct interface to the disk array and is used only by the HP service representative. |
| **TID** | Target ID. |
| **VSC** | Volume Size Configuration is a feature that defines custom volumes (CVS volumes) that are smaller than normal fixed-sized logical disk devices (volumes). |
| **WWN** | World Wide Name. A unique identifier assigned to a Fibre Channel device. |

# Index


**A**

about this guide 5
adapters
  host bus 21
arbitrated-loop physical address 18
audience
  intended 5
authorized reseller, HP 7

**C**

clustering 30, 31
command devices, RAID Manager 15
Command View XP 12
configure devices 36
configuring ports 18
connecting
  disk array 32
connection problems
  troubleshooting 44
conventions
  documentation 6

**D**

device specifications 47
device types
  supported 13
devices
  add 34
  configuring 36
  label 37
  verification 35
disk array
  connecting 32
  installation 17
disk arrays
  supported 5
documentation
  conventions 6
  related 5

**E**

emulations 46

**F**

fabric zoning 30, 31
failover 14
features and requirements 12
fiber parameter settings 19
Fibre Channel adapters
  configuring 21
  verify installation 29
Fibre Channel interface 13
Fibre Channel ports 18

file system
create 39

G
glossary 61

H
HBA driver 21
HBAs
configuring 21
help
obtaining 7
host mode 20
HP
authorized reseller 7
technical support 7

I
installation
disk array 17
interface
Fibre Channel 13

L
LUN Configuration Manager
define LUN mapping 33
LUN Configuration Manager XP 12
LUN management 12
LUN security 30, 31

O
optional software 12

P
partition
create 37
label 37
paths
defining 33
ports
configuring 18
Fibre Channel 18

R
RAID Manager command devices 15
related documentation
list of 5
Remote Control 12
requirements and features 12

S
SCSI paths
add 34
security, LUN 30, 31
software options 12
system administrator
required knowledge 5
system option mode 17

T
technical support
HP 7
troubleshooting
error conditions 42
HP Support Center 44

W
warranty 9

X
XP1024
support for 5
XP128
support for 5

XP48
support for 5
XP512
adapter types supported 13
support for 5

## Z

zoning, fabric 30, 31